arXiv:1403.6752v1 [math.ST] 26 Mar 2014

Confidence intervals for high-dimensional inverse covariance estimation

Jana Janková and Sara van de Geer

Seminar für Statistik
ETH Zürich
Switzerland

Abstract: In this paper, we establish confidence intervals for individual parameters of a sparse concentration matrix in a high-dimensional setting. We follow the idea of the projection approach proposed in [15], which is applied to the graphical Lasso to obtain a de-sparsified estimator. Subsequently, we analyze the asymptotic properties of the novel estimator, establishing asymptotic normality and confidence intervals for the case of sub-Gaussian observations. Performance of the proposed method is illustrated in a simulation study.

Keywords and phrases: covariance, concentration, sparsity, high-dimensional, graphical Lasso, inference.

1. Introduction

Motivation. A large number of methods has been proposed for the problem of inverse covariance estimation in high-dimensional settings, where the number of parameters is much larger than the sample size. For an overview of the different methodologies, we refer to [11]. However, few theoretical results exist on assessing the accuracy of these estimators and developing methodology for statistical inference for the parameters of interest in this setting.
In this paper, our aim is to develop methodology for inference for individual entries of the inverse covariance matrix. To this end, we consider a popular estimator of the inverse covariance matrix, mostly referred to as the graphical Lasso [5]. We note that existing work on statistical inference in high dimensional settings has mostly focused on inference for parameters in linear models and generalized linear models [15, 17, 8, 6, 3, 2, 10]. Our setting lies outside of the regression context, nevertheless, as shall be demonstrated, a general approach presented in [15] carries through successfully to the present context.
Problem overview and related work. Consider an i.i.d. sample $X_1, \ldots, X_n \in \mathbb{R}^p$ of size n from a zero-mean distribution with unknown covariance matrix $\Sigma^* \in \mathbb{R}^{p \times p}$. Denoting the inverse covariance matrix, often referred to as the precision or concentration matrix, as $\Theta^* = (\Sigma^*)^{-1}$, the goal is to the estimate Θ^* in a setting where $p \gg n$. The most natural candidate for an estimator of the covariance matrix is presumably the sample covariance matrix. However, if we consider the setting $p > n$, the sample covariance matrix is singular with proba-

bility one. Moreover, even when p/n tends to a constant, the covariance matrix exhibits poor performance [7]. Consequently, different structural assumptions have been imposed on the model to allow for consistent estimation. One natural assumption is a certain kind of sparsity in the model. For instance, sparsity may be imposed by restricting the cardinality of the set of non-zero elements of the covariance matrix or precision matrix.

A popular estimation procedure which relies on the sparsity in Θ^* is the graphical Lasso [5, 4, 16]. The estimator minimizes the negative Gaussian log-likelihood with regularization in terms of the ℓ_1 norm of the off-diagonal (or possibly all) entries of the precision matrix. Although the method arises from the situation when the data is normally distributed, it has justification also when this is not the case. For a discussion see [11], where it is shown that for non-Gaussian X the method corresponds to minimizing the ℓ_1 penalized log-determinant Bregman divergence. The optimization problem corresponding to graphical Lasso is a convex optimization problem that can be solved with, for instance, coordinate descent methods [4, 5] in polynomial time.

The statistical performance of graphical Lasso has been analyzed in [13] and [11], where convergence rates of the method were established under a sparsity assumption on the entries of Θ^*. In [13] (see also [9]), mild conditions on the eigenvalues of Θ^* allow to derive rates in Frobenius norm. High-dimensionality is reflected in p being allowed to grow as a function of n, however, in limit, $p/n \to 0$ is required. In [11] convergence rates for the maximum norm are derived under the stronger irrepresentability condition on the true precision matrix Θ^*, but only a weaker condition $\log p/n \to 0$ is required, allowing for $p \gg n$ asymptotically.

A closely related work [15] considers inference for low-dimensional subsets of a high-dimensional vector in the regression setting under sparsity assumptions on the parameters and the design matrix. The proposed method is based on the Lasso estimator for which the Karush-Kuhn-Tucker conditions are modified ("inverted") to obtain a de-sparsified estimator. Subsequently, the estimator is shown to have a normal limiting distribution and reach the semiparameteric efficiency bound.

In addition we mention a recent paper [12] that suggests a "piecewise" procedure for precision matrix estimation in Gaussian graphical models, under mild assumptions on Θ^*. The method restricts itself to the Gaussian model by making use of the well-known conditional distribution property of the multivariate normal distribution. This leads to the idea of using regression to estimate the precision matrix. Roughly speaking, the procedure suggests to regress every pair of variables (X^i, X^j) on the remaining $p-2$ variables using the scaled lasso regression, resulting in order $\binom{p}{2}$ individual high-dimensional regressions. The acquired partial estimators are then combined to form an estimator for Θ^*. Asymptotic normality is proved, allowing for inference results about the parameters.

1.1. Outline

In this paper we aim for statistical inference for individual elements of the precision matrix in a high-dimensional setting under appropriate sparsity assumptions on Θ^*. The work closely follows the idea of the more general approach presented in [15], which builds on "inverting" the necessary Karush-Kuhn-Tucker conditions for an optimization problem. The paper [15] demonstrates this method for the case of linear regression and generalized linear models, while we apply the idea to a fully nonlinear estimator. By inverting the KKT conditions, we obtain a de-sparsified graphical Lasso estimator and consequently we analyze its asymptotic properties. Asymptotic normality of the new estimator is proved for sub-Gaussian observations.
The rest of the paper is organized as follows. In section 1.2, we briefly introduce the model. Section 1.3 summarizes deviation inequalities for the covariance matrix for further reference throughout the paper. Section 2 contains the main results. Section 3 illustrates the theoretical results in a simulation study. Finally, section 4 contains proofs.

Notation. For a vector $x \in \mathbb{R}^d$ and $p \in (0, \infty]$ we use the notation $\|x\|_p$ to denote the $p-$norm of x in the classical sense. For a matrix $A \in \mathbb{R}^{d\times d}$ we use the notations $|||A|||_\infty = \max_i \|e_i^T A\|_1$, $|||A|||_1 = \left|||A^T\right|||_\infty$ and $\|A\|_\infty = \max_{i,j} |A_{ij}|$. The symbol $\mathrm{vec}(A)$ denotes the vectorized version of a matrix A. For two matrices A and B, use $A \otimes B$ to denote the Kronecker product of A and B. By e_i we denote a p-dimensional vector of zeros with one at position i and by $e_{ij} := e_i \otimes e_j$ a p^2-dimensional vector of zeros with one at position indexed by (i, j).
For sequences f_n, g_n, we write $f_n = O(g_n)$ if $|f_n| \leq C|g_n|$ for some $C > 0$ independent of n and all $n > C$. Analogously, we write $f_n = \Omega(g_n)$ if $|f_n| \geq C|g_n|$ for some $C > 0$ independent of n and all $n > C$. We write $f_n \asymp g_n$ if both $f_n = \mathcal{O}(g_n)$ and $f_n = \Omega(g_n)$ hold. Finally, $f_n = o(g_n)$ if $\lim_{n\to\infty} f_n/g_n = 0$.
We use S_+^p to denote the cone of positive semi-definite $p \times p$ matrices, i.e. $S_+^p := \{A \in \mathbb{R}^{p\times p} | A = A^T, A \succeq 0\}$ and S_{++}^p to denote the set of positive definite $p \times p$ matrices, $S_{++}^p := \{A \in \mathbb{R}^{p\times p} | A = A^T, A \succ 0\}$.

1.2. Model setup

In this section, we introduce the model and present assumptions that will be referenced throughout the paper. Given an i.i.d. sample $X_1, \ldots, X_n$, the graphical Lasso estimator $\hat{\Theta}_n$ is defined as the solution to the optimization problem

$$\hat{\Theta}_n := \arg\min_{\Theta\in S_{++}^p} \left\{ \mathrm{trace}(\Theta^T \hat{\Sigma}) - \log\det(\Theta) + \lambda \|\Theta\|_{1,\mathrm{off}} \right\}, \qquad \textbf{(P1)}$$

where $\hat{\Sigma}^n$ is the sample covariance matrix $\hat{\Sigma}^n = \frac{1}{n}\sum_{i=1}^n X_i X_i^T$ and $\|\cdot\|_{1,\mathrm{off}}$ is the ℓ_1 off-diagonal penalty, $\|\Theta\|_{1,\mathrm{off}} = \sum_{i\neq j} |\Theta_{ij}|$. Note that when the data is normally distributed, the problem is equivalent to ℓ_1-penalized maximum likelihood for the precision matrix.

The setting when the observations are Gaussian is perhaps of the highest interest due to the fact that the problem corresponds to estimation of the structure of an underlying Gaussian graphical model. Assume thus in the rest of this paragraph that the data comes from a multivariate normal distribution with zero mean and a covariance matrix Σ^*, i.e. $X = (X^1, \ldots, X^p) \sim \mathcal{N}(0, \Sigma^*)$. Associate the variables $X^1, \ldots, X^p$ with the vertex set $V = \{1, \ldots, p\}$ of an undirected graph $G = (V, E)$ with an edge set E. We say that the concentration matrix Θ^* respects the edge structure of the graph G if $\Theta^*_{ij} = 0$ for all $(i, j) \notin E$. Consequently, estimation of the precision matrix corresponds to parameter estimation in a Gaussian graphical model and specifying the non-zero set of Θ^* corresponds to model selection. Note that in the Gaussian case, sparsity assumptions on the entries of the precision matrix translate to sparsity of the edges in the underlying Gaussian graphical model.

Model assumptions. Let $S \equiv S(\Theta^*) := \{(i, j) \in V \times V : \Theta^*_{ij} \neq 0\}$ be the set of all non-zero entries of Θ^* and denote the cardinality of S by s. Use $S^c(\Theta^*)$ for the complement of $S(\Theta^*)$ in $V \times V$. We shall impose a sparsity assumption on the maximum row cardinality of Θ^*; therefore define $d = d_n$ as follows

$$d := \max_{i \in \{1,\ldots,p\}} |\{j \in V : \Theta^*_{ij} \neq 0\}|.$$

In terms of a Gaussian graphical model, the parameter d corresponds to the maximum vertex degree in the graph G (note that also self-loops at each vertex are counted). The parameter d is allowed to grow as a function of n.
We shall consider two parameters of the model, κ_{Σ^*} and κ_{Γ^*} defined below (cf. [11]) which might be viewed as certain measures of the model complexity. The parameters may be allowed to grow as a function of the sample size n and the rate of this growth then determines the convergence rates $\|\hat{\Theta} - \Theta^*\|_\infty$. However, to keep the presentation cleaner, we assume the parameters κ_{Σ^*} and κ_{Γ^*} are bounded in n.
First, define κ_{Σ^*} to be the ℓ_∞ operator norm of the true covariance matrix Σ^*, i.e.

$$\kappa_{\Sigma^*} = |||\Sigma^*|||_\infty = \max_i \sum_{j=1}^{p} |\Sigma^*_{ij}|.$$

To define the second parameter, consider first the Hessian $\Gamma(\Theta)$ of the function $\mathrm{tr}(\Theta^T \hat{\Sigma}) - \log\det(\Theta)$, which takes the form

$$\Gamma(\Theta) = \Theta^{-1} \otimes \Theta^{-1}.$$

We shall impose some restrictions on the Hessian Γ evaluated at the true Θ^*, $\Gamma^* := \Gamma(\Theta^*)$. To this end, let us fix the following notation. For any two subsets T and T' of $V \times V$, we use $\Gamma^*_{TT'}$ to denote the $|T| \times |T'|$ matrix with rows and columns of Γ^* indexed by T and T' respectively.
Consequently, define κ_{Γ^*} to be the ℓ_∞ operator norm of the matrix

$$\Gamma^*_{SS} = [\Theta^{*-1} \otimes \Theta^{*-1}]_{SS} \in \mathbb{R}^{s \times s}.$$

i.e., $\kappa_{\Gamma^*} = \left|\left|\left|(\Gamma^*_{SS})^{-1}\right|\right|\right|_\infty$.

Finally, we shall consider an irrepresentability condition on the Hessian of the log-determinant function evaluated at Θ^* as in assumption **(A1)** below, and boundedness of eigenvalues as in assumption **(A2)**.

Assumption (A1). *(Irrepresentability condition) There exists $\alpha \in (0, 1]$ such that*

$$\max_{e \in S^c} \|\Gamma^*_{eS}(\Gamma^*_{SS})^{-1}\|_1 \leq 1 - \alpha. \tag{1.1}$$

Remark 1. *Theoretically, under the irrepresentability condition, one could use the graphical Lasso, and then the maximum likelihood estimator based on non-zeros of $\hat{\Theta}$. The de-sparsified graphical Lasso is an alternative that we believe to be less sensitive to violations of the conditions for exact variable selection. Besides, the above procedure would result in discarding the "false negatives" produced by the graphical Lasso. We also refer here to [14] where the results of [11] are extended using an irrepresentable condition on the small (not necessarily zero) entries of Θ^*.*

Assumption (A2). *(Bounded eigenvalues) There exists $M > 0$ such that, for every n, each eigenvalue Λ of Θ^* satisfies $\frac{1}{M} < \Lambda < M$.*

Assumption (A3). *(Complexity parameters) The parameters $\kappa_{\Sigma^*}, \kappa_{\Gamma^*}$ are bounded by a universal constant for all n.*

Remark 2. *A simple consequence of assumption **(A2)** is upper-boundedness of $|\Theta^*_{ij}|$ and $|\Sigma^*_{ij}|$ by a universal constant for all i, j, and that $|\Theta^*_{ii}|$ and $|\Sigma^*_{ii}|$ are uniformly bounded away from zero.*

Since our results depend on the convergence rates for graphical Lasso derived in [11], the assumption **(A1)** corresponds to the irrepresentability condition used in [11]. Parameters appearing in assumption **(A3)** influence the rates of convergence if they scale with n (see [11]).

1.3. Tail Bounds

In this section, based on [11] and [1], we review a result concerning a high-probability bound for a deviation condition for the sample covariance matrix. Since the optimization problem **(P1)** depends on a surrogate of Σ^* which is the sample covariance matrix $\hat{\Sigma}$, the performance of the method (such as the rates $\|\hat{\Theta} - \Theta^*\|_\infty$) depends on the difference $\hat{\Sigma} - \Sigma^*$. Roughly speaking, the role of $\hat{\Sigma} - \Sigma^*$ arises by considering the estimation error $|\text{trace}\{(\hat{\Theta} - \Theta^*)(\hat{\Sigma} - \Sigma^*)\}| \leq \|\hat{\Sigma} - \Sigma^*\|_\infty \|\hat{\Theta} - \Theta^*\|_1$. The general approach (as in [11]) depends on showing that for a specific pre-assumed statistical model for X, the event

$$\mathcal{T}_n := \{\|\hat{\Sigma}^n - \Sigma^*\|_\infty < \delta_n\} \tag{1.2}$$

holds with high probability for $n \to \infty$ and some sequence $\delta_n \searrow 0$.
To measure the rate of decay of the tails of $\hat{\Sigma}^n_{ij} - \Sigma^*_{ij}$, we introduce the following definition.

Definition 1. *Given a random vector $X \in \mathbb{R}^p$, we say that a function $\tau : \mathbb{N} \times (0,\infty) \to \mathbb{R}^+$, increasing in both variables, is a tail function if there exists a constant $\nu \in (0,\infty]$ such that for all $(i,j) \in V \times V$ and for all $\delta \in (0,\nu]$,*

$$\mathbb{P}(|\hat{\Sigma}^n_{ij} - \Sigma^*_{ij}| \geq \delta) \leq \frac{1}{\tau(n,\delta)},$$

where $\hat{\Sigma}$ is the sample covariance matrix associated with an i.i.d. sample distributed as X. For a given tail function τ, define the inverse tail function for $r \geq 1$ (and for each n fixed)

$$\delta_\tau(r;n) := \arg\max\{\delta | \tau(n,\delta) \leq r\}. \tag{1.3}$$

For a given tail behaviour of X, the union bound gives a high probability bound for $\mathcal{T}_n$ as in lemma 6.

Sub-Gaussian Tails

The deviation condition $\|\hat{\Sigma}^n - \Sigma^*\|_\infty < \delta_n$ is subsequently shown to hold with high probability for special classes of random variables. Of interest is for instance the case when X is sub-Gaussian, though the results can be extended to some other cases as well (cf. [11]). We shall restrict ourselves to the class of sub-Gaussian random variables.

Definition 2. *A real zero-mean random variable X is sub-Gaussian if there exists $\beta > 0$ such that for all $t > 0$*

$$\mathbb{P}(|X| > t) \leq 2e^{-\frac{t^2}{2\beta^2}}.$$

We shall consider the following sub-Gaussianity condition for random vectors.

Condition (C1). *(Sub-Gaussianity condition) A zero-mean random vector $X = (X^1, \ldots, X^p)$ with a covariance matrix Σ^* satisfies the sub-Gaussianity condition if all its normalized components $X^i/\sqrt{\Sigma^*_{ii}}$ are sub-Gaussian random variables with a common parameter $\beta > 0$.*

The following lemma from [11] gives probabilistic bounds for the event $\mathcal{T}_n$ if X satisfies **(C1)**.

Lemma 1 ([11], sub-Gaussian model)**.** *Let $X = (X^1, \ldots, X^p)$ be a zero mean random vector with covariance matrix Σ^* satisfying the sub-Gaussianity condition **(C1)**. Then an inverse tail function is given by*

$$\delta_\tau(r,n) = c_0\sqrt{\frac{\log(4r)}{n}},$$

and for every $\gamma > 2$ and each n such that $\delta_\tau(n,p^\gamma) < c_1$ we have

$$\mathbb{P}\left(\|\hat{\Sigma}^n - \Sigma^*\|_\infty \geq \delta_\tau(n,p^\gamma)\right) \leq \frac{1}{p^{\gamma-2}},$$

where c_0, c_1 depend only on β and $\max_i \Sigma^*_{ii}$.

2. Main results

In this section we present the main results which imply inference for individual parameters of the precision matrix. The main result is contained in theorem 1, where asymptotic normality of the de-sparsified estimator is proved. At this point we remark that for any $\lambda_n > 0$ and $\hat{\Sigma}^n$ with strictly positive diagonal elements, the optimization problem **(P1)** has a unique solution $\hat{\Theta}_n \in S^p_{++}$ which is characterized by the Karush-Kuhn-Tucker conditions

$$\hat{\Sigma} - \hat{\Theta}^{-1} + \lambda \hat{Z} = 0, \tag{2.1}$$

where the matrix $\hat{Z}$ belongs to the subdifferential of the off-diagonal norm $\|\cdot\|_{1,\text{off}}$ evaluated at $\hat{\Theta}$ (see Lemma 3 in [11]).
Inverting the KKT conditions (2.1) leads us to the de-sparsified graphical Lasso estimator which is defined as follows

$$\hat{T} := \text{vec}(\hat{\Theta}) - \hat{\Theta} \otimes \hat{\Theta} \text{vec}(\hat{\Sigma} - \hat{\Theta}^{-1}). \tag{2.2}$$

To prove asymptotic normality of $\hat{T}$, we shall first need the following lemma, which shows that a term we may identify as a "remainder term" converges in probability to zero. The result holds uniformly for all the elements of Θ^*. The proof of the lemma is given in section 4.

Lemma 2. *Suppose that Θ^* satisfies assumptions **(A1)**,**(A3)** and the sparsity assumption $d = \mathcal{O}(\frac{1}{\delta_n})$ for some $0 < \delta_n \leq 1$. Suppose that $\hat{\Theta}_n$ is the solution to the optimization problem **(P1)** with sample covariance matrix $\hat{\Sigma}$ and regularization parameter $\lambda_n = \frac{8}{\alpha}\delta_n$. Then*

$$\sqrt{n}(\hat{T} - vec(\Theta^*)) + W = Rem, \tag{2.3}$$

where $W = \sqrt{n}\Theta^* \otimes \Theta^* vec(\hat{\Sigma} - \Sigma^*)$ and on the set $\mathcal{T}_n = \{\|\hat{\Sigma} - \Sigma^*\| < \delta_n\}$, it holds that

$$\|Rem\|_\infty \lesssim \sqrt{n} \max\{d^{\frac{3}{2}}\delta_n^2, d^2\delta_n^3\}.$$

The statement of lemma 2 is deterministic, but as lemma 1 suggests, in cases of interest (e.g. the sub-Gaussian model), the deviation condition holds with high probability.
The result of lemma 2 implies that under an additional sparsity assumption $\sqrt{n}\max\{d^{\frac{3}{2}}\delta_n^2, d^2\delta_n^3\} = o(1)$, on the event $\mathcal{T}_n$ it holds that $\|Rem\|_\infty = o(1)$.

Corollary 1 (Sub-Gaussian model)**.** *Suppose that in addition to the assumptions of lemma 2, the i.i.d. random vectors $X_1, \ldots, X_n \in \mathbb{R}^p$ satisfy the sub-Gaussianity condition **(C1)**. Put $\delta_n \asymp \delta_\tau(n, p^\gamma) = o(1)$, where $\gamma > 2$. Then it follows that*

$$\|Rem\|_\infty = \mathcal{O}_P\left(d^{\frac{3}{2}}\frac{\log p}{\sqrt{n}}\right).$$

Furthermore, provided that $d^{\frac{3}{2}} = o(\frac{\sqrt{n}}{\log p})$ it follows that $\|Rem\|_\infty = o_P(1)$.

Note also that if $X_1, \ldots, X_n \in \mathbb{R}^p$ is an i.i.d. sample from the multivariate normal $\mathcal{N}_p(0, \Sigma^*)$ distribution, then W in lemma 2 has the Wishart distribution $\mathcal{W}_p(\Theta^*, n)$ (up to scaling and shifted by its mean).

Consequently, lemma 2 under the sub-Gaussian model assumption lets us show asymptotic normality as demonstrated in theorem 1 below. Since our aim is inference about individual elements of Θ^*, fix $(i, j) \in V \times V$ and to keep the notation simple, we suppress the index (i, j). Note also that the columns of Θ^* will be denoted by $\Theta^*_{\cdot i} \in \mathbb{R}^p$.
For a fixed n the elements of the sequence of random vectors $X_1, X_2, \ldots, X_n$ depend also on $p = p_n$ ($X_k \in \mathbb{R}^p$). In order to control this dependence we shall consider the doubly indexed (or triangular) array $\{X_{nk}\}_{1 \leq k \leq n}, n \in \mathbb{N}$.

Theorem 1. *Let $\{X_{nk}\}_{k=1}^n$ be i.i.d. satisfying the sub-Gaussianity condition **(C1)** with parameter β which does not depend on n. Suppose that $\hat{\Theta}_n$ is the solution to the optimization problem **(P1)** with regularization parameter $\lambda_n = \frac{8}{\alpha}\delta_\tau(n, p^\gamma)$ for some $\gamma > 2$. Let Θ^* satisfy assumptions **(A1)**,**(A2)**,**(A3)** and the sparsity assumption $d^{\frac{3}{2}} = o(\frac{\sqrt{n}}{\log p})$. Suppose that*

$$\sigma_n^2 := Var({\Theta^*_{\cdot i}}^T X_{n1} {\Theta^*_{\cdot j}}^T X_{n1}) = \Omega(1).$$

Then for all $(i, j) \in V \times V$, it holds that

$$\sqrt{n}(\hat{T}_{ij} - \Theta^*_{ij})/\sigma_n = Z^n_{ij} + o_P(1), \tag{2.4}$$

where Z^n_{ij} converges weakly to $\mathcal{N}(0, 1)$.

Note that in practice, σ_n is unknown, so instead of σ_n one needs to use a consistent estimator $\hat{\sigma}_n > 0$. The next lemma shows that for the Gaussian model $\sigma_n^2 = \Omega(1)$ and there is a natural estimator of σ_n^2.

Lemma 3. *Suppose that the conditions of Theorem 1 hold with the exception of $\sigma_n^2 = \Omega(1)$. Furthermore, suppose that X_{nk} are Gaussian random vectors. Then for all $(i, j) \in V \times V$, it holds that*

$$\sqrt{n}(\hat{T}_{ij} - \Theta^*_{ij})/\hat{\sigma}_n = Z^n_{ij} + o_P(1),$$

where $\hat{\sigma}_n^2 := \hat{\Theta}^2_{ij} + \hat{\Theta}_{ii}\hat{\Theta}_{jj}$, Z^n_{ij} converges weakly to $\mathcal{N}(0, 1)$.

3. Simulation Results

In this section we report on the performance of our method on simulated data and provide a brief comparison to some other methodologies. We consider two particular instances of sparse Gaussian graphical models which may be (fully) specified by the precision matrix Θ^*.
The model setup for the simulations was specified as follows. As a first graphical model instance, we consider the chain graph on p vertices where the maximum

vertex degree is by definition restricted to $d = 2$. As a second instance, we consider the star graph where a given central vertex may or may not be connected to any of the remaining vertices. Thus the maximum degree d may vary from 0 to $p - 1$. We specify the graphical models with the corresponding precision matrices in a normalized form with ones on the main diagonal. For the case of a chain graph, the corresponding precision matrix Θ^* is a tridiagonal matrix, which was set to $\Theta^* = \text{tridiag}(\rho, 1, \rho)$ for a given $\rho > 0$. The cardinality of the active set is then $3p - 2$. For the star graph, $\Theta^*_{ii} = 1$ and $\Theta^*_{ij} = \rho$ for $i \neq j, (i, j) \in E$.

To solve the graphical Lasso program **(P1)**, we have used the implemented procedure *glasso* of Friedman et al. [5].

By corollary 3, it follows that the $(1 - \alpha)100\%$ asymptotic confidence interval for Θ^*_{ij} is given by

$$I_{ij} \equiv I_{ij}(\hat{\Theta}_{ij}, \alpha, n) := [\hat{T}_{ij} - \Phi^{-1}(1 - \alpha/2)\frac{\hat{\sigma}_{ij}}{\sqrt{n}}, \hat{T}_{ij} + \Phi^{-1}(1 - \alpha/2)\frac{\hat{\sigma}_{ij}}{\sqrt{n}}].$$

For each parameter Θ^*_{ij}, the probability that the true value Θ^*_{ij} is covered by the confidence interval was estimated by its empirical version, $\hat{\alpha}_{ij} := \mathbb{P}_N \mathbf{1}_{\{\Theta^*_{ij} \in I_{ij,\alpha}\}}$. The number of iterations used to calculate the estimates $\hat{\alpha}_{ij}$ was set to $N = 300$. Next for a set $A \subset V \times V$ define the average coverage over the set A as

$$\text{Avgcov}_A := \frac{1}{|A|} \sum_{(i,j) \in A} \hat{\alpha}_{ij}.$$

After obtaining the estimates $\hat{\alpha}_{ij}$, we have averaged them over the sets S_0 and S_0^c to obtain $Avgcov_{S_0}$ and $Avgcov_{S_0^c}$, respectively. Similarly, we have calculated the average length of the confidence interval for each parameter Θ^*_{ij} from $N = 300$ iterations and again averaged these over the sets S_0 and S_0^c to obtain $Avglength_{S_0}$ and $Avglength_{S_0^c}$. We have carried out a comparison of the graphical Lasso-based confidence intervals with two other methods for their construction: the "oracle" maximum likelihood estimator with the zero set S_0^c pre-specified and a method based on the sample covariance matrix. The results are shown in figure 1. It may be noted that the graphical Lasso performs worse for the chain graph, where the active set is bigger.

The choice of the regularization parameter. Theory implies that the correct choice of the regularization parameter satisfies $\lambda_n \asymp \sqrt{\frac{\log p}{n}}$. In our setting, roughly speaking, it might be sufficient to choose the regularization parameter λ smaller than when model selection is the goal (i.e. identifying the zero set of Θ^*). Considering confidence intervals, there is presumably not a significant difference between a small absolute value of a parameter and a zero value, thus less regularization may suffice.

The practical choice of regularization was done empirically, optimizing over average coverage as a function of λ. Very roughly, based on our experiments for the two Gaussian graphical models, $\lambda_n = \sqrt{\frac{\log p}{n}}$ produces reasonably satisfactory results.

Chain graph	S_0 Avgcov	S_0 Avglength	S_0^c Avgcov	S_0^c Avglength
graphical Lasso	0.934	0.247	0.972	0.215
MLE with specified S_0	0.963	0.293	-	-
Sample covariance	0.459	0.428	0.897	0.367

Star graph, $d = 8$	S_0 Avgcov	S_0 Avglength	S_0^c Avgcov	S_0^c Avglength
graphical Lasso	0.948	0.328	0.951	0.247
MLE with specified S_0	0.956	0.337	-	-
Sample covariance	0.124	0.499	0.897	0.367

FIG 1. *Tables showing a comparison of graphical Lasso with the maximum likelihood estimator with specified set S_0 and an estimator based on the sample covariance matrix $\hat{\Sigma}$. The results were obtained for $p = 80$ and $n = 250$. The regularization parameter was chosen $\lambda = \sqrt{\frac{\log p}{n}}$. The active sets have cardinality 238 for the chain graph and cardinality 94 for the star graph. The constant ρ in the definition of Θ^* was set to 0.3 for both graphical models.*

Chain graph	S_0 Avgcov	S_0 Avglength	S_0^c Avgcov	S_0^c Avglength
p = 100	0.931	0.401	0.978	0.348
p = 200	0.917	0.400	0.984	0.349
p = 300	0.893	0.401	0.988	0.349
p = 500	0.832	0.401	0.988	0.350

FIG 2. *A table showing the performance of the graphical Lasso* ***(P1)*** *for number of parameters p taking values $100, 200, 300, 500$ and $n = 100$. The regularization parameter was chosen $\lambda = \sqrt{\frac{\log p}{n}}$. The constant ρ is 0.3 in the definition of Θ^*.*

4. Proofs

Proof of Lemma 2. The KKT conditions for the optimization problem **(P1)** read

$$\hat{\Sigma} - \hat{\Theta}^{-1} + \lambda \hat{Z} = 0, \tag{4.1}$$

where the matrix $\hat{Z}$ is the subdifferential of $\|.\|_{1,\text{off}}$ at the optimum $\hat{\Theta}$. Vectorizing the equation and multiplying by the matrix $\hat{\Theta} \otimes \hat{\Theta}$, we obtain

$$\hat{\Theta} \otimes \hat{\Theta}\text{vec}(\hat{\Sigma}) - \hat{\Theta} \otimes \hat{\Theta}\text{vec}(\hat{\Theta}^{-1}) + \hat{\Theta} \otimes \hat{\Theta}\text{vec}(\lambda \hat{Z}) = 0.$$

Adding the expression $\hat{\Theta} - \Theta^*$ to both sides and rearranging to obtain the asymptotic pivot (i.e. subtracting expectation from $\hat{\Sigma}$), it follows

$$\begin{aligned} \text{vec}(\hat{\Theta}) + \hat{\Theta} \otimes \hat{\Theta}\text{vec}(\lambda \hat{Z}) - \text{vec}(\Theta^*) &= -\Theta^* \otimes \Theta^* \text{vec}(\hat{\Sigma} - \Sigma^*) \\ &\quad + \underbrace{(\Theta^* \otimes \Theta^* - \hat{\Theta} \otimes \hat{\Theta})\text{vec}(\hat{\Sigma} - \Sigma^*)}_{Rem_1} \\ + \underbrace{\text{vec}(\hat{\Theta} - \Theta^*) - \hat{\Theta} \otimes \hat{\Theta}\text{vec}(\Theta^{*-1}) + \hat{\Theta} \otimes \hat{\Theta}\text{vec}(\hat{\Theta}^{-1})}_{Rem_2}. \end{aligned} \tag{4.2}$$

Expressing $\lambda\hat{Z}$ from equation (4.1), on the LHS of (4.2) we obtain the expression $\hat{T} - \text{vec}(\Theta^*)$.
Next we find an elementwise bound for the remainders Rem_1 and Rem_2. In what follows, condition on the event $\mathcal{T}$ defined in (1.2).

Denoting $\Delta = \hat{\Theta} - \Theta^*$ and rewriting the second remainder in terms of Δ we get

$$\begin{aligned} \|Rem_2\|_\infty = \|\Delta{\Theta^*}^{-1}\Delta\|_\infty &= \max_{i,j\in\{1,\dots,p\}} |e_i^T\Delta{\Theta^*}^{-1}\Delta e_j| \\ &\leq \max_i \max_j \|e_i\Delta\|_\infty \|{\Theta^*}^{-1}\Delta e_j\|_1 \qquad (4.3) \\ &= \|\Delta\|_\infty \left|\!\left|\!\left|\Delta{\Theta^*}^{-1}\right|\!\right|\!\right|_\infty. \qquad (4.4) \end{aligned}$$

where inequality (4.3) follows since $|a^Tb| \leq \|a\|_\infty\|b\|_1$ for any vectors $a, b \in \mathbb{R}^p$ and equality (4.4) follows since $\max_i \|e_iA\|_\infty = \|A\|_\infty$ and $\max_j \|Ae_j\|_1 = |||A|||_1 = \left|\!\left|\!\left|A^T\right|\!\right|\!\right|_\infty$.

Next by the sub-multiplicativity of $|||.|||_\infty$ it follows that

$$\|\Delta{\Theta^*}^{-1}\Delta\|_\infty \leq \|\Delta\|_\infty |||\Delta|||_\infty \left|\!\left|\!\left|{\Theta^*}^{-1}\right|\!\right|\!\right|_\infty. \qquad (4.5)$$

By Lemma 5, part (a), on $\mathcal{T}$ it holds that $\hat{\Theta}_{S_0^c} = \Theta^*_{S_0^c}$. Thus $\hat{\Theta}$ has at most d nonzero entries per row. Hence it follows

$$|||\Delta|||_\infty := \left|\!\left|\!\left|\hat{\Theta} - \Theta^*\right|\!\right|\!\right|_\infty \leq d\|\hat{\Theta} - \Theta^*\|_\infty. \qquad (4.6)$$

Continuing with (4.5) we obtain by lemma 5, part (b),

$$\|\Delta{\Theta^*}^{-1}\Delta\|_\infty \leq d\kappa_{\Sigma^*}\|\Delta\|_\infty^2 \leq d\kappa_{\Sigma^*}\delta_n^2.$$

Next we bound the second remainder $\|Rem_1\|_\infty = \|(\hat{\Theta}\otimes\hat{\Theta} - \Theta^*\otimes\Theta^*)\text{vec}(\hat{\Sigma} - \Sigma^*)\|_\infty$.
For any $m \in \mathbb{N}$, matrix $A \in \mathbb{R}^{m\times m}$ and a vector $x \in \mathbb{R}^m$ we have $\|Ax\|_\infty \leq |||A|||_\infty\|x\|_\infty$. Thus it follows that

$$\begin{aligned} \|Rem_1\|_\infty &= \|(\hat{\Theta}\otimes\hat{\Theta} - \Theta^*\otimes\Theta^*)\text{vec}(\hat{\Sigma} - \Sigma^*)\|_\infty \\ &\leq \left|\!\left|\!\left|\hat{\Theta}\otimes\hat{\Theta} - \Theta^*\otimes\Theta^*\right|\!\right|\!\right|_\infty \|\hat{\Sigma} - \Sigma^*\|_\infty. \end{aligned}$$

Recall that conditioned on $\mathcal{T}$, we have $\|\hat{\Sigma} - \Sigma^*\|_\infty \leq \delta_n$. Also observe that $\hat{\Theta}\otimes\hat{\Theta} - \Theta^*\otimes\Theta^*$ may be rewritten as $\Theta^*\otimes\Delta + \Delta\otimes\Theta^* + \Delta\otimes\Delta$. Hence

$$\left|\!\left|\!\left|\hat{\Theta}\otimes\hat{\Theta} - \Theta^*\otimes\Theta^*\right|\!\right|\!\right|_\infty \|\hat{\Sigma} - \Sigma^*\|_\infty \leq |||\Theta^*\otimes\Delta + \Delta\otimes\Theta^* + \Delta\otimes\Delta|||_\infty\delta_n$$

For any matrices $A, B \in \mathbb{R}^{m\times m}$ it holds that $|||A\otimes B|||_\infty = |||A|||_\infty|||B|||_\infty$. To see this, realize that the $(i,j)-$th row of the Kronecker product is given by $(A\otimes$

$B)_{(i,j),\cdot} = A_{i\cdot} \otimes B_{j\cdot}$. For vectors u, v, it is easy to check that $\|u \otimes v\|_1 = \|u\|_1 \|v\|_1$ by straightforward calculation. Hence $|||A \otimes B|||_\infty = \max_{(i,j)\in V\times V} \|(A\otimes B)_{(i,j),\cdot}\|_1$ $= \max_{(i,j)\in V\times V} \|A_{i\cdot}\|_1 \|B_{j\cdot}\|_1 = \max_{i\in V} \|A_{i\cdot}\|_1 \max_{j\in V} \|B_{j\cdot}\|_1 = |||A|||_\infty |||B|||_\infty$. Applying the triangle inequality together with the last property of $|||\cdot|||_\infty$ gives

$$\begin{aligned} |||\Theta^* \otimes \Delta + \Delta \otimes \Theta^* + \Delta \otimes \Delta|||_\infty \delta_n &\leq \{2|||\Delta \otimes \Theta^*|||_\infty + |||\Delta \otimes \Delta|||_\infty\} \delta_n \\ &= \left\{2|||\Delta|||_\infty |||\Theta^*|||_\infty + |||\Delta|||_\infty^2\right\} \delta_n \end{aligned}$$

Applying (4.6) again and noting that $|||\Theta^*|||_\infty = \max_i \|\Theta_{i\cdot}\|_1 \leq \sqrt{d} \max_i \|\Theta_{\cdot i}\|_2 \leq \sqrt{d}\Lambda_{\max}(\Theta^*)$, we obtain

$$\begin{aligned} \left\{2|||\Delta|||_\infty |||\Theta^*|||_\infty + |||\Delta|||_\infty^2\right\} \delta_n &\leq \left\{2\kappa_{\Sigma^*} d^{\frac{3}{2}} \Lambda_{\max}(\Theta^*)\delta_n + d^2\delta_n^2\right\} \delta_n \\ &\leq 2\max\{2\kappa_{\Sigma^*}\Lambda_{\max}(\Theta^*) d^{\frac{3}{2}}\delta_n^2, d^2\delta_n^3\}. \end{aligned}$$

Finally, we obtain the following elementwise bound for the remainder

$$\begin{aligned} \|Rem\|_\infty &\leq \|Rem_1\|_\infty + \|Rem_2\|_\infty \\ &\leq 2\sqrt{n}\max\{\kappa_{\Sigma^*} d\delta_n^2, 4\kappa_{\Sigma^*}\Lambda_{\max}(\Theta^*) d^{\frac{3}{2}}\delta_n^2, 2d^2\delta_n^3\} \\ &\lesssim \sqrt{n}\max\{d^{\frac{3}{2}}\delta_n^2, d^2\delta_n^3\}. \end{aligned}$$

where the last inequality follows since by assumption **(A3)** κ_{Σ^*} is uniformly bounded, by assumption **(A2)** $\Lambda_{\max}(\Theta^*) < M$, and since $d \geq 1$. □

Proof of theorem 1. By corollary 1, under the sub-Gaussianity condition **(C1)**, the assumptions of lemma 2 and the sparsity assumption $d^{\frac{3}{2}} = o(\frac{\sqrt{n}}{\log p})$, it holds that $\|Rem\|_\infty = o_P(1)$. Considering then the result of lemma 2 elementwise, for every (i, j) it holds that

$$\begin{aligned} \sqrt{n}(\hat{T}_{ij} - \Theta^*_{ij}) &= \sqrt{n}\sum_{k=1}^n \frac{1}{n}(e_{ij}^T(\Theta^* \otimes \Theta^*)\mathrm{vec}(X_{nk}X_{nk}^T) - \Theta^*_{ij}) + o_P(1), \\ &= \frac{1}{\sqrt{n}}\sum_{k=1}^n (\Theta^{*T}_{\cdot i} X_{nk} \Theta^{*T}_{\cdot j} X_{nk} - \Theta^*_{ij}) + o_P(1). \end{aligned} \tag{4.7}$$

To prove the claim of the theorem, we need to show that the scaled summation term in (4.7) weakly converges to the normal distribution. To this end, define

$$Z_{nk} := \Theta^{*T}_{\cdot i} X_{nk} \Theta^{*T}_{\cdot j} X_{nk} - \Theta^*_{ij}.$$

For each n fixed, $Z_{n1}, \ldots, Z_{nn}$ are identically distributed r.v.s with mean zero since $\mathbb{E}(Z_{nk}) = \Theta^*_{\cdot i}\Sigma^*\Theta^*_{\cdot j} - \Theta_{ij} = e_i^T\Theta^*\Sigma^*\Theta^* e_j - \Theta^*_{ij} = 0$.
To show that $\sigma_n^2 \equiv \sigma_{nk}^2 := Var(Z_{nk}) = Var(\Theta^{*T}_{\cdot i} X_{nk} \Theta^{*T}_{\cdot j} X_{nk})$ is finite, note that for each n fixed, $\Theta^{*T}_{\cdot i} X_{nk}$ is sub-Gaussian since a finite sum of sub-Gaussian random variables is sub-Gaussian. The moments of a sub-Gaussian random variable are finite, which implies finiteness of σ_n^2 since $\mathbb{E}[(\Theta^{*T}_{\cdot i} X_{nk} \Theta^{*T}_{\cdot j} X_{nk})^2] \leq$

$\mathbb{E}[(\Theta^{*}_{\cdot i}{}^{T} X_{nk})^4] + \mathbb{E}[(\Theta^{*}_{\cdot j}{}^{T} X_{nk})^4]$. Denote the partial sum of Z_{nk}'s by $S_n := \sum_{k=1}^{n} Z_{nk}$ and its variance $s_n^2 := Var(S_n)$. Then $s_n^2 = n\sigma_n^2$. Dividing (4.7) by $\sigma_n > 0$, we obtain

$$\sqrt{n}(\hat{T}_{ij} - \Theta^{*}_{ij})/\sigma_n = S_n/s_n + o_P(1)/\sigma_n.$$

To prove the claim of the theorem we need to show that $S_n/s_n \xrightarrow{D} \mathcal{N}(0,1)$ and $o_P(1)/\sigma_n = o_P(1)$. The latter statement follows by the assumption $\sigma_n = \Omega(1)$. If the Lindeberg condition is satisfied, i.e. for all $\varepsilon > 0$

$$\lim_{n\to\infty} \frac{1}{s_n^2} \sum_{j=1}^{n} \mathbb{E}(Z_{nj}^2 \mathbf{1}(|Z_{nj}| > \varepsilon s_n)) = 0,$$

then we have $\frac{S_n}{s_n} \xrightarrow{D} \mathcal{N}(0,1)$ for $n \to \infty$.
Observe that since $\{Z_{nk}\}_{k=1}^{n}$ are identically distributed for n fixed and $s_n^2 = n\sigma_n^2$,

$$\sum_{k=1}^{n} \frac{1}{s_n^2} \mathbb{E}(Z_{nk}^2 \mathbf{1}(|Z_{nk}| > \varepsilon s_n)) = \frac{1}{\sigma_n^2} \mathbb{E}(Z_{n1}^2 \mathbf{1}(|Z_{n1}| > \varepsilon s_n)).$$

Consequently, it remains to show that $\lim_{n\to\infty} \mathbb{E}(Z_{n1}^2 \mathbf{1}(|Z_{n1}| > \varepsilon s_n)) = 0$. For any $c > 0$ we may rewrite

$$c = \int_0^{\infty} \mathbf{1}_{c>t} dt.$$

Therefore, applying the last observation and by Fubini's theorem we obtain

$$\int_{\Omega} |X|^2 \mathbf{1}_{|X|>a} d\mathbb{P} \quad = \quad a^2 \mathbb{P}(|X| > a) + 2 \int_a^{\infty} u \mathbb{P}(|X| > u) du.$$

Then it follows

$$\mathbb{E}(Z_{n1}^2 \mathbf{1}(|Z_{n1}| > \varepsilon\sigma_n\sqrt{n})) \leq \varepsilon^2\sigma_n^2 n \mathbb{P}(|Z_{n1}| > \varepsilon\sigma_n\sqrt{n}) + \int_{\varepsilon\sigma_n\sqrt{n}}^{\infty} x \mathbb{P}(|Z_{n1}| > x) dx. \tag{4.8}$$

To show that the limit of the right-hand side of the last inequality is 0 for $n \to \infty$, we use lemma 4, by which it follows that Z_{n1} satisfies a tail bound $\mathbb{P}(|Z_{n1}| > t) \leq 4de^{-\frac{\lambda}{c_1 d}}$. For a fixed $\varepsilon > 0$, putting $t := \varepsilon\sigma_n\sqrt{n}$, we obtain

$$\mathbb{P}(|Z_{n1}| > \varepsilon\sigma_n\sqrt{n}) \leq 4de^{-\frac{\varepsilon\sigma_n\sqrt{n}}{c_1 d}}.$$

Using the last inequality, we show that the limit of the last two terms in (4.8) for $n \to \infty$ is 0. For the first term we have

$$\lim_{n\to\infty} \sigma_n^2 n P(|Z_{n1}| > \varepsilon\sigma_n\sqrt{n}) \leq \lim_{n\to\infty} \sigma_n^2 n d e^{-\frac{\varepsilon\sigma_n\sqrt{n}}{c_1 d}} = 0, \tag{4.9}$$

which follows by assumption $d^{\frac{3}{2}} = o(\frac{\sqrt{n}}{\log p})$. Next considering the limit of the last term in (4.8), we have

$$\int_{\varepsilon\sigma_n\sqrt{n}}^{\infty} xP(|Z_{n1}| > x)dx \leq \int_{\varepsilon\sigma_n\sqrt{n}}^{\infty} 4xde^{-\frac{x}{c_1 d}}dx.$$

In the integral, substitute $t := \frac{x}{\sigma_n\sqrt{n}}$ to obtain

$$\lim_{n\to\infty} \int_{\sigma_n\sqrt{n}}^{\infty} dxe^{-\frac{x}{d}}dx \quad = \quad \lim_{n\to\infty} \int_{1}^{\infty} d\sigma_n^2 nte^{-\frac{\sigma_n\sqrt{n}}{d}t}dt.$$

Again by the restriction on d and the Lebesgue dominated convergence it then follows that the limit of the integral is 0. In conclusion, we get $\frac{S_n}{s_n} \xrightarrow{D} \mathcal{N}(0,1)$ for $n \to \infty$. □

Lemma 4. *Let Θ^* satisfy assumption **(A2)** and let the random vector $X_n \in \mathbb{R}^p$ satisfy the sub-Gaussianity condition **(C1)**. Then for $t > c_0$ the random variable $Z_n := {\Theta^*_{\cdot i}}^T X_n {\Theta^*_{\cdot j}}^T X_n - \Theta^*_{ij}$ satisfies the following bound*

$$\mathbb{P}(|Z_n| > t) \leq 4de^{-\frac{t}{c_1 d}},$$

where c_0, c_1 do not depend on n.

Proof. Since $|{\Theta^*_{\cdot i}}^T X_n| \leq \|\Theta_{\cdot i}\|_1 \max_{k:\Theta^*_{ki}\neq 0} |X_n^k| \leq \|\Theta^*_{\cdot i}\|_2\sqrt{d}\max_{k:\Theta^*_{ki}\neq 0}|X_n^k|$, and using the union bound and sub-Gaussianity of $X_n^k/\sqrt{\Sigma^*_{kk}}$ by assumption **(C1)** gives

$$\begin{aligned}
\mathbb{P}(|{\Theta^*_{\cdot i}}^T X_n| > t) &\leq \mathbb{P}\left(\max_{k:\Theta^*_{ki}\neq 0} |X_n^k| > \frac{t}{\sqrt{d}\|\Theta^*_{\cdot i}\|_2}\right) \\
&\leq d \max_{k:\Theta^*_{ki}\neq 0} \mathbb{P}\left(|X_n^k|/\sqrt{\Sigma^*_{kk}} > \frac{t}{\sqrt{d}\|\Theta^*_{\cdot i}\|_2\sqrt{\Sigma^*_{kk}}}\right) \\
&\leq 2d\exp\left(-\frac{t^2}{2\beta^2 d\|\Theta^*_{\cdot i}\|_2^2 \max_k \Sigma^*_{kk}}\right).
\end{aligned}$$

Note that by remark 2, Σ^*_{kk} and $\|\Theta^*_{\cdot i}\|_2 \leq \Lambda_{\max}(\Theta^*)$ are uniformly bounded in n by the assumption on eigenvalues of Σ^*.
Then for ${\Theta^*_{\cdot i}}^T X_n {\Theta^*_{\cdot j}}^T X_n$ we obtain the bound

$$\mathbb{P}(|{\Theta^*_{\cdot i}}^T X_n {\Theta^*_{\cdot j}}^T X_n| > t) \leq 4d\exp\left(-\frac{t}{2\beta^2 d\Lambda^2_{\max}(\Theta^*)\max_k \Sigma^*_{kk}}\right).$$

Thus for $Z_n = {\Theta^*_{\cdot i}}^T X_n {\Theta^*_{\cdot j}}^T X_n - \Theta^*_{ij}$, there exist constants c_0, c_1, c_2 not depending on n such that for $t > c_0 > |\Theta^*_{ij}|$

$$\mathbb{P}(|Z_n| > t) \leq 4de^{-\frac{t-|\Theta^*_{ij}|}{c_2 d}} \leq 4de^{-\frac{t}{c_1 d}},$$

since $|\Theta^*_{ij}|/d$ is bounded by **(A2)** and $d \geq 1$. □

Proof of lemma 3. The claim of the corollary follows by first applying Theorem 1 provided that we show that $\sigma_n = \Omega(1)$ and then by multiplying (2.4) by $\sigma_n/\hat{\sigma}_n$ provided that we show $\sigma_n/\hat{\sigma}_n \xrightarrow{P} 1$.
We shall first prove that σ_n^2 is the corresponding element of the inverse of Fisher information matrix for Θ^*. Assuming that $X \sim \mathcal{N}(0, \Sigma^*)$, then ${\Theta^*}^T X \sim \mathcal{N}(0, \Theta^*)$, hence

$$\sigma_n^2 = \mathrm{Var}({\Theta^*_{\cdot i}}^T X {\Theta^*_{\cdot j}}^T X) = \mathrm{Var}(e_i^T {\Theta^*}^T X {\Theta^*}^T X e_j) = \mathrm{Var}(e_i^T Z Z^T e_j),$$

where $Z \sim \mathcal{N}(0, \Theta^*)$. Thus

$$\sigma_n^2 = Var(Z_i Z_j) = \mathbb{E}(Z_i^2 Z_j^2) - \mathbb{E}(Z_i Z_j)^2 = \Theta^*_{ii}\Theta^*_{jj} + 2{\Theta^*_{ij}}^2 - {\Theta^*_{ij}}^2 = \Theta^*_{ii}\Theta^*_{jj} + {\Theta^*_{ij}}^2.$$

Next we show that $\sigma_n = \Omega(1)$. To obtain a lower bound on the diagonal elements of Θ^*, again observe by assumption **(A2)** that $\Theta^*_{ii} \geq \Lambda_{\min}(\Theta^*) > \frac{1}{M} > 0$. Hence $\Theta^*_{ii}\Theta^*_{jj} + {\Theta^*}^2_{ij} \geq \frac{1}{M^2} > 0$, where M is independent of n, therefore $\sigma_n = \Omega(1)$ as required.
Next we show $\sigma_n/\hat{\sigma}_n \xrightarrow{P} 1$. Observe that on $\mathcal{T}$

$$\begin{aligned} |\hat{\sigma}_n^2 - \sigma_n^2| &\leq |\hat{\Theta}_{ii}\hat{\Theta}_{jj} - \Theta^*_{ii}\Theta^*_{jj}| + |\hat{\Theta}_{ij}^2 - {\Theta^*}^2_{ij}| \\ &\leq |\Delta_{ii}\Delta_{jj} + \Theta^*_{ii}\Delta_{jj} + \Theta^*_{jj}\Delta_{ii}| \\ &\quad + |\Delta_{ij}(\Delta_{ij} + 2\Theta^*_{ij})| \\ &\lesssim \max\{|\Theta^*_{jj}|, |\Theta^*_{ii}|, |\Theta^*_{ij}|, 1\}\sqrt{\frac{\log p}{n}}. \end{aligned} \tag{4.10}$$

Again by assumption **(A2)**, the absolute values $|\Theta^*_{jj}|, |\Theta^*_{ii}|, |\Theta^*_{ij}|$ are bounded, which together with (4.10) implies (since $\mathcal{T}$ holds with high probability) that $\lim_{n\to\infty} \mathbb{P}(|\hat{\sigma}_n - \sigma_n| > \varepsilon) = 0$ for all $\varepsilon > 0$. This in turn implies that $\hat{\sigma}_n/\sigma_n \xrightarrow{P} 1$ since $\sigma_n = \Omega(1)$. But then also $\sigma_n/\hat{\sigma}_n \xrightarrow{P} 1$.
Finally, note that $o_P(1)\sigma_n/\hat{\sigma}_n = o_P(1)$ to prove the claim.
As a final remark, note that the expression $\hat{\Theta}_{ii}\hat{\Theta}_{jj} + \hat{\Theta}_{ij}^2 \geq \lambda^2_{\min}(\hat{\Theta}) > 0$ is strictly positive since $\hat{\Theta}$ is positive definite by definition **(P1)** and therefore we may divide by it for every n. □

References

[1] Bickel, P. J. and Levina, E. (2008). Regularized estimation of large covariance matrices. *Annals of Statistics* **36** 199-227.

[2] Chatterjee, A. and Lahiri, S. N. (2011). Bootstrapping Lasso estimators. *Journal of the American Statistical Association* **106** 608-625.

[3] Chatterjee, A. and Lahiri, S. N. (2013). Rates of convergence of the adaptive LASSO estimators to the oracle distribution and higher order refinements by the bootstrap. *ArXiv:1307.1952*.

[4] d'Aspremont, A., Banerjee, O. and El Ghaoui, L. (2008). First-order methods for sparse covariance selection. *SIAM J. Matrix Anal. Appl.* **30** 56–66.

[5] Friedman, J., Hastie, T. and Tibshirani, R. (2008). Sparse inverse covariance estimation with the graphical lasso. *Biostatistics* **9** 432–441.
[6] Javanmard, A. and Montanari, A. (2013). Confidence intervals and hypothesis testing for high-dimensional regression. *ArXiv: 1306.3171*.
[7] Johnstone, I. M. (2001). On the distribution of the largest eigenvalue in principal components analysis. *The Annals of Statistics* **29** 295–327.
[8] Meinshausen, N. (2013). Assumption-free confidence intervals for groups of variables in sparse high-dimensional regression. *ArXiv:1309.3489*.
[9] Negahban, S. N., Ravikumar, P., Wainwright, M. J. and Yu, B. (2010). A unified framework for high-dimensional analysis of M-estimators with decomposable regularizers. *Statistical Science* **27** 538-557.
[10] Nickl, R. and van de Geer, S. (2012). Confidence sets in sparse regression. *The Annals of Statistics* **41** 2852-2876.
[11] Ravikumar, P., Raskutti, G., Wainwright, M. J. and Yu, B. (2008). High-dimensional covariance estimation by minimizing l1-penalized log-determinant divergence. *Electronic Journal of Statistics* **5** 935-980.
[12] Ren, Z., Sun, T., Zhang, C. H. and Zhou, H. H. (2013). Asymptotic normality and optimalities in estimation of large Gaussian graphical model. *ArXiv: 1309.6024*.
[13] Rothman, A. J., Bickel, P. J., Levina, E. and Zhu, J. (2008). Sparse permutation invariant covariance estimation. *Electronic Journal of Statistics* **2** 494–515.
[14] van de Geer, S. (2014). Worst possible sub-directions in high-dimensional models. *Working paper*.
[15] van de Geer, S., Bühlmann, P., Ritov, Y. and Dezeure, R. (2013). On asymptotically optimal confidence regions and tests for high-dimensional models. *ArXiv:1303.0518*.
[16] Yuan, M. and Lin, Y. (2007). Model selection and estimation in the Gaussian graphical model. *Biometrika* 117.
[17] Zhang, C. H. and Zhang, S. S. (2014). Confidence intervals for low-dimensional parameters in high-dimensional linear models. *Journal of the Royal Statistical Society: Series B* **76** 217-242.

Appendix A

In this section, we restate a result on rates of convergence of the graphical Lasso from Ravikumar et al. [11] in lemma 5 and a high-probability bound for $\mathcal{T}_n$ in lemma 6.

Lemma 5 (Theorem 1, Ravikumar et al. [11])**.** *Let $\hat{\Theta}_n$ be the solution to the optimization problem **(P1)**. Let the assumption **(A1)** hold, put $\lambda_n = \frac{8}{\alpha}\delta_n$ for some $\delta_n > 0$ and let the sparsity assumption $\delta_n \leq \frac{1}{c_1 d}$ hold. Then on $\mathcal{T}_n$,*

(a) $\hat{\Theta}_{S_0^c} = \Theta^*_{S_0^c}$
(b) $\|\hat{\Theta} - \Theta^*\|_\infty \leq c_2\delta_n$,

where c_1, c_2 depend only on $\kappa_{\Sigma^*}, \kappa_{\Gamma^*}$ and α.

Statement (a) is a consequence of the result that, under the assumptions of lemma 5, the graphical Lasso estimator $\hat{\Theta}$ is equal to the "oracle" penalized maximum likelihood estimator

$$\tilde{\Theta}_{oracle} := \arg \min_{\Theta \in \mathcal{S}^p_{++}, \Theta_{S_0^c} = \Theta^*_{S_0^c}} \left\{ \text{trace}(\Theta^T \hat{\Sigma}) - \log \det(\Theta) + \lambda \|\Theta\|_{1,\text{off}} \right\} \quad \textbf{(P2)}$$

on $\mathcal{T}$ (see Lemma 4 in Ravikumar et al. [11]).

For instance, given sub-Gaussian observations, lemma 5 in combination with lemma 1 gives the following convergence rates. If the sparsity assumption $d = \mathcal{O}(\sqrt{\frac{n}{\log p}})$ is satisfied (equivalently $n = \Omega(d^2 \log p)$) and for suitably chosen $\lambda_n \asymp \delta_n$ we have

$$\|\hat{\Theta} - \Theta^*\|_\infty = \mathcal{O}_P \left(\sqrt{\frac{\log p}{n}} \right).$$

Lemma 6 (Lemma 8, Ravikumar et al. [11])**.** *Let X be a zero-mean random vector with precision matrix Θ^*, inverse tail function δ_τ and a tail constant ν. For $\gamma > 2$ put $\delta_n := \delta_\tau(n, p^\gamma)$ in the definition of $\mathcal{T}_n$. Then for a sample size n such that $\delta_\tau(n, p^\gamma) < \nu$, we have*

$$P(\mathcal{T}_n) \geq 1 - \frac{1}{p^{\gamma - 2}}.$$